Institute for Scientific Computing Research

ISCR Subcontract Research Summaries

Large Scalable Algebraic Domain Decomposition Preconditioners

Randolph E. Bank
University of California, San Diego

Summary:

Many leading-edge scientific and engineering simulations expend large amounts of computational resources for the solution of linear systems of equations. Multilevel and domain decomposition methods have been identified as potentially scalable linear system solvers on terascale computer platforms. However, interprocessor communication overheads, degree of parallelism (in solving the coarse problems), and effects of increasing number of processors on convergence rates all contribute to the list of obstacles to true scalability. The novel feature of our domain decomposition approach is that the subproblem residing in each processor is defined over the entire domain, although the vast majority of unknowns for each subproblem are associated with the subdomain owned by the corresponding processor. The purpose is to ensure that a global coarse description of the problem is contained within each of the subproblems, so a separate coarse grid solve is not required to achieve rapid convergence of the overall iteration. This procedure is similar in philosophy to the parallel adaptive mesh refinement paradigm introduced by Bank and Holst, in that it attempts to both maximize the use of existing sequential algebraic solvers on each processor, and minimize as much as possible the communication between processors. This approach can be applied to general sparse matrices, although matrices arising from discretizations of partial differential equations are the principal target.

In a classic domain decomposition (DD) algorithm, each processor is responsible for solving a linear system corresponding to its subdomain. The DD preconditioner can often be improved by adding overlap, typically allowing processors to solve for certain unknowns near the interface but outside its domain. Additionally, a coarse grid correction step is often applied to capture the global low frequency error modes, a necessary step for rapid convergence. Our philosophy extends these ideas to a natural conclusion; in some sense each processor is given the entire problem to solve, so there is maximal overlap, allowing each processor to independently resolve the global low frequency errors. However, on the parts of the problem it does not own, it uses a coarsened approximation. These subproblems are solved locally on each processor, using any efficient sequential solver and no communication. Our hope is that these solves comprise the majority of work involved in the overall solution process, and the number of outer iterations will be small due to rapid convergence of the global iteration, leading to an efficient scalable algorithm with low communication.

A preliminary version of a DD procedure based on this paradigm was implemented in the finite element package PLTMG, and the computational results are quite promising. It appears from two to ten global iterations are necessary for the class of problems solved by PLTMG (scalar elliptic PDEs), with each global iteration requiring just two communication steps — a boundary exchange near the beginning of each iteration, and a few scalars (norms and inner products used in convergence criteria and damping) near the bottom of the iteration. However, in PLTMG the issue of coarsening does not arise directly. The problems which are refined in a given processor's subdomain and coarse in the remaining regions are provided naturally by the Bank–Holst parallel adaptive meshing paradigm.

Summary (continued):

Randolph E. Bank
University of California, San Diego

Another algorithm based on this approach was proposed by Bank and Jimack and analyzed in a manuscript by Bank, Jimack, Nadeem, and Nepomnyaschikh, where it was shown to be optimal. Here again, the structure of the coarse mesh was provided a priori, and the details of this grid structure played an important role in the theoretical analysis.

The next and major step in the project will be to develop a similar algorithm, basing the implementation as much as possible on existing routines in the *hypre* library at LLNL. Since *hypre* already has excellent algebraic solvers that can be used in this application, the main focus will be on building the coarse approximation for regions outside of the given processor's subdomain. This coarsening problem is related to the construction of algebraic multilevel pre-conditioners, but it is simplified because a hierarchy of levels is not needed. It is hoped that existing routines in *hypre* can be adapted to the requirements of this approach. If successful, this inherently parallel DD method will offer an attractive and easily implemented alternative for solving large sparse systems of equations arising from discretized PDEs.

Multilevel Nonlinear Additive Schwarz Preconditioned Inexact Newton Methods and Applications

Xiao-Chuan Cai
University of Colorado

Summary:

The focus of our research has been on the design, analysis, and software implementation of a new family of multilevel additive Schwarz preconditioned inexact Newton methods for solving nonlinear algebraic systems of equations arising from the finite element or finite difference discretization of nonlinear partial differential equations. Inexact Newton is a very fast algorithm for solving systems whose nonlinearities are well balanced; however, such balance rarely exists in the solution of many practically important problems, such as high Reynolds number incompressible flows or transonic compressible flows. A technique of interest for such cases is to precondition the nonlinear system before calling the inexact Newton solver. This idea has recently been proven effective when using a single-level additive Schwarz method such as the nonlinear preconditioner (ASPIN).

The single level ASPIN is partly scalable in the sense that when the mesh is refined the number of nonlinear outer Newton iterations doesn't change much, and the number of global linear iterations for solving the Jacobian system doesn't change much either. However, for elliptically dominated problems when the number of processors increases, though the number of Newton iterations remains fixed, the number of global linear iterations per Newton iteration increases considerably. This limits the usefulness of the method on ASCI-scale computers. To make ASPIN scalable with respect to the number of processors in a parallel computer, we are currently developing several multilevel versions.

One approach is based on the classical two-level additive Schwarz method in which a coarse solution is simply added to the local solutions. We have showed theoretically that the method provides a nonlinear system that is equivalent to the original system and showed numerically that the number of global linear iterations does not increase much when the number of processors increases.

Our other approach is based on the FAS method, which uses the coarse space correction in a multiplicative way. The complexity inherent in such software for parallel multilevel nested nonlinear iterations demands a computational framework that provides broad-based infrastructure. Our software is written using Argonne National Lab's PETSc, which integrates a hierarchy of components that range from low-level distributed data structures for grids, vectors, and matrices through high-level linear and nonlinear solvers.

Several test cases have been studied using ASPIN. This includes the incompressible Navier–Stokes equation, the compressible full potential equation, and a non-resistive magnetohydrodynamic flow problem.

First-Order System Least Squares for Linear and Nonlinear Elasticity Equations

Zhiqiang Cai
Purdue University

Summary:

Basic equations of elasticity are generally in self-adjoint form, so they lend themselves naturally to an energy minimization principle, cast in terms of the primitive displacement variables. Unfortunately, this direct approach seems to have many practical difficulties (e.g., degraded approximation properties of the discretization and convergence properties of the solution process) as the material tends to become incompressible (i.e., the Poisson ratio tends to 0.5 from below). There have been several attempts to develop alternative approaches that are robust in the incompressible limit. Compounding these difficulties is the fact that what is often needed in practice is the elements of the stress tensor. These variables can be obtained by differentiating displacements, but this weakens the order and strength of the approximation.

The practical need of the stress tensor motivated extensive studies of mixed-finite element methods in the stress-displacement formulation. Unlike mixed methods for second-order scalar elliptic boundary value problems, stress-displacement finite elements are extremely difficult to construct. This is because the stress tensor is symmetric. A beautiful finite element space had not been constructed until recently by Arnold and Winther. Their space is a natural extension of the Raviart–Thomas space of H(div). Previous work imposed the symmetry condition weakly via a Lagrange multiplier. The minimum degree of freedom on each triangle of Arnold–Winther space for the symmetric stress tensor in two dimensions is twenty-four, which is very expensive. Like scalar elliptic problems, mixed methods lead to saddle-point problems and mixed finite elements are subject to the inf-sup condition. Many solution methods that work well for symmetric positive problems cannot be applied directly. Although substantial progress in solution methods for saddle-point problems has been achieved, these problems may still be difficult and expensive to solve.

In the recent years there has been a serious interest in least-squares methods. A number of least-squares formulations have been proposed, analyzed, and implemented. In particular, the least-squares method by Cai, Manteuffel, and McCormick aims to compute the stress tensor directly and, hence, accurately, and it is robust in the incompressible limit. This method is a two-stage algorithm that first solves for the gradients of displacement (which immediately yield stress tensor), then for the displacement itself (if desired). Under certain H2 regularity assumptions, it admits optimal H1-like performance for standard finite element discretization and standard multigrid and domain decomposition solution methods that is uniform in the Poisson ratio for all variables. A limitation of this approach is the requirement of sufficient smoothness of the original problem. Also, the gradient of displacement is not an immediate physical quantity and it is hard to extend this approach to nonlinear elasticity.

With goals of the accurate approximation to the stress, robustness in the incompressible limit, efficient solvers, and applicability to nonlinear elasticity, we developed a least-squares finite element method based on the stress-displacement formulation. As we mentioned before, a major numerical difficulty is how to handle the symmetry of the stress tensor in the stress-displacement formulation. To circumvent such a difficulty, we impose the symmetry condi-

Summary (continued):

Zhiqiang Cai
Purdue University

tion in the first-order system and then apply the least-squares principle to this over-determined, but consistent system. The least-squares functional uses the L2 norm and it is shown that the homogeneous functional is equivalent to the energy norm involving the Lamé constant for the displacement and the standard H(div) norm for the stress. This implies that our least-squares finite element method, using the respective Crouzeix–Raviart and Raviart–Thomas spaces for the displacement and stress, yields optimal error estimates uniform in the incompressible limit. The total number of degrees of freedom is twelve per triangle in two dimensions and eighteen per tetrahedron in three dimensions. This work has been written up for submission to *SIAM J. Numer. Anal.* The algebraic system resulting in this discretization may be efficiently solved by multigrid methods, which is the topic of our current study. We will also continue our effort by the important extension of this approach to nonlinear elasticity and possible applications in the ALE3D project in LLNL.

Investigation of the Richtmyer–Mechkov Instability Using Adaptive Mesh and Algorithm Refinement

Alejandro Garcia
San Jose State University

Summary:

In recent years there has been a growing interest in fluid mechanics at microscopic length and time scales. At these scales the continuum representation of hydrodynamics, as expressed by partial differential equations, breaks down and computational fluid mechanics (CFD), whose numerical schemes are based on such a macroscopic description, is not accurate. For such regimes, algorithms based on statistical mechanics have been developed and shown to capture accurately all the relevant physics. The most well-known examples of such simulation methods are molecular dynamics (MD) and direct simulation Monte Carlo (DSMC); the greatest disadvantage of these molecular simulations is their computational expense relative to CFD methods.

For some problems only a small (but important) volume in a calculation requires molecular resolution. This scenario is reminiscent of adaptive mesh refinement where only a limited region, say near a shock front, requires high resolution. This association led to the development of Adaptive Mesh and Algorithm Refinement (AMAR), which embeds a DSMC simulation within an adaptive mesh refinement framework.

Recently, Rich Hornung (LLNL), Sanith Wijesinghe (MIT), and I have developed an advanced AMAR code using SAMRAI (Structured Adaptive Mesh Refinement Applications Infrastructure). We have also benefited from the assistance of SAMRAI team members (Steve Smith, Andy Wissink, et al.) and Nicholas Hadjiconstantinou, Wijesinghe's doctoral advisor at MIT. Using the SAMRAI framework has allowed us to design a code that has multiple DSMC regions and is fully adaptive; multiple refinement and de-refinement criteria can be imposed. SAMRAI-based communications routines handle the passing of particle information between patches (e.g., copying particles from one DSMC patch into the ghost cells of another). The hierarchical data in the simulation is output in the form of plot files formatted for the Vizamrai suite of visualization tools.

The first application for which the AMAR/SAMRAI code was originally targeted is the Richtmyer–Meshkov instability, that is, the enhancement of mixing in a binary gas due to the impulsive acceleration from a shock wave. This past year multi-species functionality was added to the Euler solver. A variety of test cases have been designed to test these new routines and their coupling with the DSMC portion of the program. Until recently it was difficult to obtain quantitative results from the Vizamrai-based output files, but this summer Steve Smith provided us with new data reduction routines to "crunch" these hierarchical data files into a simplified form that can be processed by standard numerical tools such as Matlab. The analysis of this data, as well as further validation testing, is currently underway.

Despite the improved efficiency of a hybrid code versus a conventional particle simulation, our problems of interest require large-scale computations best performed on a massively parallel computer. Our work this summer showed that the heavy communications and memory demands of a particle method require the development of more sophisticated load balancing strategies. The potentially high cost of particle data redistribution during the adaptive gridding phases of the solution algorithm should also be investigated. Recently the memory management in the DSMC routines was redesigned, which we hope will allow us to run significantly larger problems. Parallel algorithm questions such as these are of increasing interest in computational science research, especially for multi-scale and multi-algorithm methods such as AMAR.

Lagrange–Newton–Krylov–Schur Solvers for PDE-Constrained Optimization

Omar Ghattas
Carnegie Mellon University

Summary:

This project focuses on Lagrange–Newton–Krylov–Schur (LNKS) algorithms for large-scale optimization problems that are constrained by systems of partial differential equations and their applications to optimal control, optimal design, and parameter estimation problems in science and engineering.

Collectively, we refer to these as "inverse problems" to distinguish them from "forward problems" that usually characterize large-scale simulations. In the forward problem, problem data—initial conditions, boundary conditions, material coefficients, and the domain geometry—are specified, and the state of the system is found by solving the PDEs. The inverse problem involves the reverse process: some components of the state are typically specified (through an objective function to be minimized), and solution of the PDE-constrained optimization problem yields components of the data, often called the decision variables.

The inverse problem is often significantly more difficult to solve than the forward problem. This is because the inverse problem includes the PDEs as part of its constraint set, and because the inverse problem is often ill-posed despite the well-posedness of the forward problem. Not surprisingly, most of the work in large-scale simulation has been directed at the forward problem. Sustained advances over the past twenty years have produced a body of efficient parallel scalable algorithms for many classes of PDE simulations. This invites research into what is often the ultimate goal in many areas of computational science and engineering: the optimal design, optimal control, or parameter estimation problem, in the form of a PDE-constrained optimization problem.

During my sabbatical at LLNL, I worked on LNKS algorithms for time-dependent PDE-constrained optimization in conjunction with CMU graduate student Volkan Akcelik. One of the difficulties with time-dependent problems is that optimization "sees" the entire time history, and therefore the optimization problem is coupled across time. We tailored Newton–Krylov-type methods for time domain inverse wave propagation problems, which serve as a good testbed for time-dependent optimization. The specific class of problems studied was estimation of material parameters in an acoustic medium, given waveform observations at specified locations on the boundary. The Newton–Krylov optimization methods were implemented on top of the parallel PDE solver library PETSc. Problems of up to several hundred thousand inversion parameters were solved, on up to 128 T3E processors. The mesh-independence of the Newton and Krylov iterations demonstrated algorithmic scalability with respect to number of inversion and state parameters. We also studied issues related to total variation regularization, and to overcoming the multiple minima problem via multiscale extension of the Newton–Krylov solver.

In addition to conducting research on algorithms for time-dependent inversion, during my stay at LLNL I taught a ten-lecture short course entitled "Computational Optimization." The focus of the course was on numerical algorithms for large-scale continuous optimization problems and their particularization for optimization problems that are constrained by discretized

Summary (continued):

Omar Ghattas
Carnegie Mellon University

PDEs. Sensitivity analysis and its relationship to reduced space optimization methods were also discussed. Course participants came from CASC and several science and engineering divisions.

During my stay I also participated in the Nonlinear Solvers and Differential Equations project, and I interacted with the sensitivity analysis group. CASC collaborator Carol Woodward and I worked on formulating optimal control and inversion problems for two problems in environmental and atmospheric sciences that are of interest to LLNL researchers.

A Fast Poisson Solver for General 3D Regions

Anne Greenbaum
University of Washington

Summary:

We are developing a software package for solving Poisson's equation on general three-dimensional regions, in approximately twice the time required for a fast Poisson solver on an embedding cube. The surface of the region must be triangulated, but the interior need never be meshed. The algorithm involves two steps, the first of which consists of evaluating a volume integral and the second of which requires solving Laplace's equation with given boundary conditions. Each can be expressed as a Poisson problem on an embedding cube, provided the discontinuities in the desired solution and its derivatives across the boundary of the region can be determined. This requires solving a second Fredholm boundary integral equation. Such equations are very amenable to solution by iterative methods such as the GMRES algorithm. If there are m boundary discretization points, then the work to solve the integral equation using the GMRES algorithm with dense matrix-vector multiplication is O(m**2), but this can be reduced to O(m) work if the matrix-vector multiplications are carried out with the Fast Multipole Method. The remaining work is that of two Poisson solves on a regular lattice throughout an embedding cube. If there are N lattice points, this is O(N log N) work using an FFT and can even be carried out with O(N) work using multigrid methods (though the constant may be larger). We are writing the code in a modular way so that different algorithms can be substituted for the individual pieces to see which performs best.

The algorithm is amenable to parallelization, and we have tested parts of it on an IBM SP2. The goal was to solve problems on a lattice of (2N)**3 points using 8p processors in about the same amount of time required to solve a problem of size N**3 on p processors—i.e., to have a scalable algorithm. With the fast Poisson solver, the total work is O(N log N) instead of O(N), so we did not expect perfect scalability, but since log N is such a small factor, we did see near perfect scalability. This was very encouraging and shows that the algorithm parallelizes efficiently. We plan to incorporate this package in the LLNL Djehuty code for stellar evolution. Another advantage of this approach, which will be of particular use in the Djehuty code, is that derivatives can be computed directly. This avoids numerical difficulties associated with computing the solution and then using a finite difference formula for its derivatives. Other advantages of this approach include the fact that it can handle exterior problems as easily as interior ones. As mentioned earlier, one never meshes the domain but only the boundary. Additionally, some other equations such as the biharmonic equation, which might prove more difficult for finite difference or finite element techniques, can be solved with this approach almost as easily as Poisson's equation.

Some Aspects on Algebraic Multilevel and Multigrid Methods

Johannes K. Kraus
University of Leoben, Austria

Summary:

Algebraic multilevel and multigrid methods have gained rising popularity in the solution of large sparse systems of linear equations. This is because they often combine the efficiency of classical (structured) multigrid methods with the applicability to unstructured-grid problems. In many situations algebraic multilevel preconditioners, constructed in the framework of incomplete LU decomposition, provide an alternative to classical algebraic multigrid (AMG) methods.

In the first phase of this research project at CASC/LLNL, which was supported by the Austrian Science Foundation (FWF), AMG methods were further examined. The main focus was on a comparative analysis of linear and nonlinear algebraic multilevel iteration (AMLI) based on special hierarchical ordering strategies for the unknowns. In both cases the same assumption on the approximation, caused by neglecting fill-in terms during the LU decomposition, was used. Theoretical and experimental investigations showed that the resulting variable-step preconditioners become very close to linear mappings if we employ approximations to the Schur complements that are close enough on all levels of cyclic reduction. In this case very few, e.g., two or three, inner GCG-type iterations on certain levels, e.g., every other level, are sufficient to achieve optimal order of computational complexity. It is also possible to show similar convergence rates as for the conjugate gradient method accelerated by linear AMLI. Compared to the W-cycle variant of linear AMLI the nonlinear method has the advantage of being free of any method parameters to be estimated. The main results of this work are documented in the paper "An Algebraic Preconditioning Method for M-matrices: Linear Versus Nonlinear Multilevel Iteration," which has been accepted in *Linear Algebra with Applications*.

Another research goal in this project was to investigate element-free interpolation in element-based algebraic multigrid (AMGe) methods. In joint work with Van Emden Henson and Panayot Vassilevski, a new algorithm for constructing neighborhood matrices to be used for the computation of interpolation weights was developed. The method utilizes the existence of simple interpolation matrices (piecewise constant for example) on a hierarchy of coarse spaces (grids). Then one constructs by algebraic means graded away coarse spaces for any given fine-grid neighborhood. Next, the corresponding stiffness matrix is computed on this graded away mesh, and the actual neighborhood matrix is simply the Schur complement of this matrix where degrees of freedom outside the neighborhood have been eliminated. Once the neighborhood matrices have been computed they can serve as "element matrices" in AMGe methods. A description of the algorithm with model complexity analysis as well as some comparative tests of the quality of the resulting improved interpolation matrices is to be submitted for publication ("Computing Interpolation Weights in AMG Based on Multilevel Schur Complements").

Further research was done on new multilevel incomplete factorization algorithms that result in better convergence properties when applied to matrices that are not M-matrices.

Non-Conforming Finite Elements, Mesh Generation, Adaptivity, and Related Solution Methods

Raytcho Lazarov and Joseph Pasciak

Texas A&M University

Summary:

We studied the problem of developing algebraic multilevel algorithms in a parallel computing environment. Specifically, we started with a partitioning of the original domain into subdomains with a generally unstructured finite element mesh on each subdomain. The meshes do not need to be aligned across the subdomain boundaries. An element agglomeration-based algebraic multilevel coarsening is then applied independently in each subdomain. Note that even if one starts with a conforming line grid, independent coarsening generally leads to non-matching grids on the coarser levels. The element-based coarsening gives rise to a face decomposition of each subdomain boundary such that every coarse face is a union of fine faces. We developed a general dual basis mortar approach to set up global problems on all levels based on the above coarsening. The resulting "algebraic" mortar approach is well-defined and consistent in the sense that the generalized dual basis functions reproduce constants locally.

As usual, continuity across the interfaces is imposed by integral constraints against the mortar functions and reduces to a standard dual basis approach when applied to conventional finite element spaces, e.g., the fine level discretization. A general code was developed to illustrate the behavior of the proposed method. This code is based on a number of public domain tools and tools under current development at LLNL. One starts with a conforming coarse grid which is then partitioned into sub-domains using METIS, and the subdomain meshes are refined independently in parallel. An element based agglomeration procedure (AMGe) is then applied to the fine grid problems locally on the subdomains. This AMGe technique operates on and produces (on the coarser levels) generalized "elements" defined in terms of relation tables between elements—faces and degrees of freedom and the local element matrices. This information carries over to the interfaces and enables the local face based construction of the mortar.

Towards Bridging the Processor/ Memory Performance Gap for SMP Systems

Sally A. McKee
University of Utah

Frank Mueller
North Carolina State University

Summary:

Processor speeds are increasing much faster than memory speeds, and this disparity prevents many applications from making effective use of the tremendous computing power of modern microprocessors. Current access times of 50 cycles or more often cause memory performance to dominate application run time, and the processor/memory performance gap continues to grow. The problem is acute for uniprocessor machines, and even worse for symmetric multiprocessors sharing memory resources.

Together with members of CASC, we are developing tools and measures to better understand the memory performance of large, scientific applications. We are then leveraging these to design more efficient memory systems that implement Dynamic Access Optimization (DAO) techniques, which attempt to make better use of the memory system by changing the order and/or apparent location of memory references. This work is in collaboration with SRC Computers, Inc., whose reconfigurable SRC-6 SMP machine serves as the experimental vehicle for our investigations of DAO techniques.

We began by targeting streamed computations with strided access patterns. The principle of locality has guided the design of many key architectural features, including cache hierarchies and TLBs, and quantitative measures of spatial and temporal locality of reference have been useful for predicting the performance of memory hierarchy components. Unfortunately, the concept of locality is constrained to capturing memory access patterns characterized by proximity, while sophisticated memory systems are capable of exploiting other predictable access patterns. For instance, the vectors used in streamed applications lack temporal and often spatial locality, and thus have poor cache behavior. Nonetheless, their access patterns have the advantage of being predictable, and this predictability can be exploited to improve the efficiency of the memory subsystem (e.g., by reordering accesses to avoid bank conflicts or by prefetching stream data within the memory controller).

We address this in part by defining the concepts of spatial and temporal regularity, and by introducing a measure of spatial access regularity to quantify some of the predictability in access patterns. We developed an efficient, online algorithm to dynamically determine the spatial access regularity in an application's memory references, and we have demonstrated its use on a set of regular and irregular codes. We found that the use of our algorithm, with its associated overhead of trace generation, slows typical applications by at least an order of magnitude less than traditional, full-trace generation approaches. Our approach can be applied to the characterization of program access patterns and in the implementation of sophisticated, software-assisted prefetching mechanisms, and its inherently parallel nature makes it well suited for use with multi-threaded programs. We have incorporated this technology into a dynamic instrumentation framework so that we can selectively generate partial data traces through dynamic binary rewriting. This part of our work has produced a workshop paper, a poster presentation, and a conference paper submission.

We are also developing memory microbenchmarks that use hardware performance monitors to measure and categorize latencies in SMP machines. We

Summary (continued):

Sally A. McKee
University of Utah

Frank Mueller
North Carolina State University

have used these benchmarks to evaluate the memory performance of two of the LLNL machines, the results of which were presented at the IBM SP Scientific Computing User Group Meeting (SCICOMP 4).

Current work expands on these efforts and investigates memory-controller architectures and the accompanying software to implement Dynamic Access Optimizations that exploit regularity and other kinds of predictability to deliver better memory system performance.

Discretization of Neutron Transport and Automatic Parallelization

Christoph Pflaum
University of Wurzburg

Summary:

My research at CASC was directed towards two projects. The first project was a discretization of the neutron transport equation. A general problem in the discretization of neutron transport is the approximation of the function space on the sphere. It is important to find an approximation that minimizes so-called "ray effects" and that leads to a simple stiffness matrix.

In cooperation with Britton Chang and Peter Brown, we studied the finite element discretization of a function space on the sphere. The corresponding finite element space consists of piecewise constant functions. Different types of tessellation of the sphere were studied. To obtain a discretization with small ray effects, it is necessary to do a very accurate integration of the finite element functions. Numerical results show that the resulting discretization leads to smaller ray effects than the ordinate direction approximation of the sphere. Furthermore, the resulting stiffness matrix can be evaluated in a very efficient way.

The second project was on automatic parallelization with expression templates on semi-unstructured grids. Manual parallelization of a numerical algorithm tends to be difficult and time consuming. Therefore, one is interested in concepts that lead to an automatic parallelization of algorithms.

An important problem of automatic parallelization is the description of algorithms in a suitable, expressive high-level language. In case of numerical algorithms for PDEs, one would like to implement algorithms in a language that is close to the mathematical language. This can be obtained by expression templates in C++. This concept was used in the library EXPDE for the finite element approximation of PDEs. We implemented a parallel version of this library. The new parallel version of EXPDE contains a parallel grid generator for general domains in 3D, multigrid operators, and several other concepts. The automatic parallelization concept of EXPDE was applied to several PDEs, including the Stokes equations, transport equations, and the equations of elasticity. In cooperation with Rob Falgout, we developed a parallelization concept based on octrees. An important property of this parallelization concept is that the Gauss-Seidel iteration can efficiently be parallelized even in case of relatively small grids. Numerical results on ASCI Blue-Pacific showed the scalability of the parallelization concept for more than 500 processors.

It is planned to apply a finite element approximation of functions on the sphere in the ADRA project. An extension of EXPDE an adaptive grids is also planned.

AMG, Spectral AMGe, and FOSPACK

John Ruge
Front Range Scientific Computations, Inc.

Summary:

Our project focused on AMG, Spectral AMGe, and FOSPACK. Algebraic Multigrid (AMG) is actually a class of methods designed to solve a given discrete problem (generally arising from the discretization of partial differential equations) using multigrid principles by automatically choosing coarser grids, grid transfer operators, and coarser grid operators needed in multigrid cycling. In the most general case, this is done using only the information contained in the matrix itself. This has been shown to be very robust for elliptic problems and can be applied on unstructured meshes and problems with varying or even discontinuous coefficients.

Part of this project involved the examination of fill-in on coarser level matrices while using "classical" AMG codes. While overall complexity appears bounded, stencil size on coarser levels can grow, especially in large 3D problems. This can be an impediment to effective parallelization of the method. Studies on model problems showed that this effect is not due to irregular coarsening near boundaries, as had been hypothesized, since it also appeared when periodic boundary conditions were used on uniform hexahedral meshes. Some growth seems inevitable, even with uniform coarsening patterns. Interpolation truncation was only marginally effective in controlling this growth. Future work here will focus on sparser coarse grids coupled with "long range" interpolation as a treatment for this problem.

Spectral AMGe is an offshoot of classical AMG in which interpolation is constructed in an effort to more directly approximate the lower part of the spectrum of the fine grid operator, since these are the components not effectively reduced by relaxation. This approach, applicable only to finite element discretizations, requires the local stiffness matrices corresponding to the fine grid elements. An agglomeration procedure is used to partition the mesh into groups of elements, and the local matrices corresponding to each group are assembled. The eigenvectors corresponding to the small eigenvalues of these matrices are used to form the columns of interpolation matrix. Such an approach is potentially much more robust than standard AMG approaches. Tests were performed on a FOSLS formulation (see below) of 2D Helmholtz problems, verifying that the method performs well, even in the presence of large near-null spaces of the operator.

The majority of work in this project was devoted to the further development and testing of the FOSPACK code, focusing on problems arising in the ALE3D package. First-order system least squares (FOSLS) is a relatively recent approach to problem formulation, in which a PDE or system of PDEs is reformulated as a system of first-order equations, and the problem is posed as the minimization of the sum of the square of the L2 norms of the corresponding residuals. For many problems, formulations have been obtained in which it can be shown that this functional is equivalent to the square of the H1 norm of the error. One very nice consequence of this is that the resulting system has elliptic diagonal blocks that dominate the off-diagonal blocks. This makes AMG a very attractive solver for such problems. FOSPACK is a program that combines easy problem specification, assembly, and solution of FOSLS formulations. We applied the code to elasticity problems from ALE3D, demonstrating O(h**2) convergence in both displacements and stress, and began development of a suite of diagnostic tools to help the user in the formulation of the FOSLS functional in linear and nonlinear problems.

Development of Object-Oriented Tools for the Numerical Solution of Reactive Flow

Donald W. Schwendeman
Rensselaer Polytechnic Institute

Summary:

Our work involved the development of object-oriented application software, within the Overture framework of codes, for the numerical simulation of high-speed reactive flow. The mathematical model on which the software is based is the reactive Euler equations. The numerical implementation of this model is fairly general and allows for multiple reacting species and reaction rates and a general equation of state with the aim of being able to simulate experimentally observed phenomena in gas or solid explosives. The software is part of the OverBlown package of fluids codes, developed and maintained by Bill Henshaw and the Overture team at CASC. It uses overlapping grids in order to handle general domains and the A++/P++ array class library (developed by Dan Quinlan and the Overture team), which allows parallel processing. The software includes a patch grid-type adaptive mesh refinement (AMR) scheme. The software has been carefully tested for accuracy using existing codes written previously by the author, and it has been used to model the evolution to detonation of reactive samples subject to various initial conditions and within various confinement geometries.

The main work involved the continued development of software for the numerical solution of the reactive Euler equations. The equations describe the nonlinear convection of mixture-averaged hydrodynamic variables, such as density, velocity, pressure and temperature, and the convection and production of a set of reacting species that describe the mixture. The chemistry is modeled by prescribed reaction rates for each component and these rates are highly state sensitive for the problems of interest. An equation of state for the mixture is assumed and involves the heats of reaction for each of the reacting species. For a particular simulation, the equations are to be solved on a given domain subject to various initial conditions and boundary conditions.

A numerical method of solution was implemented for the reactive Euler equations. Within the Overture framework, the numerical method discretizes the equations on a set of overlapping, structured grids that cover the domain of interest. The software package Ogen generates the overlapping grid and provides geometric mapping information for each curvilinear component grid and information concerning the communication of the solution between grids in the overlap region. This information was considered to be given and thus the main task was to implement a numerical method for an individual component grid and then let the existing software, OverBlown, handle the surrounding numerical details (such as interpolation between component grids, the application of boundary conditions, time stepping, graphical interface, and more).

The numerical method chosen was a shock-capturing, Godunov-type scheme. The main new element to the software is the implementation of an AMR scheme. This is a patch-based scheme. Multiple levels of finer AMR grids are built for each base grid and are used to resolve the fine-scale structure near the reaction zone. The AMR grids communicate at the boundaries with their coarser parent grid, or with sister grids at the same level. They may also communicate across the overlap. The refinement of the grid is based on an esti-

Summary (continued):

Donald W. Schwendeman
Rensselaer Polytechnic Institute

mate of spatial and temporal gradients of the numerical solution. This is done by monitoring the magnitude of first and second differences of the solution and the magnitude of the reaction rate. The user has considerable freedom to specify and tune the AMR parameters for a particular problem.

For certain problems it is desirable to use the grid generator to introduce front-tracking grids on the base level in order to reduce the number of AMR grids needed to track and resolve a detonation wave. It may also be desirable from an accuracy point of view to confine a detonation wave to a single component grid at the base level in order to avoid the propagation of a detonation wave along a grid overlap. In order to test this numerical approach, we have modified the software to allow grid insertion and deletion. At present this can be done manually, but work is ongoing to test various numerical strategies for managing the tracking grids automatically based on the solution behavior. For example, it is possible to determine the approximate location of a detonation wave based on the behavior of the reacting species.

A number of problems involving high-speed reactive flow have been identified and are under current investigation using the new software. These problems involve paths to detonation of reactive samples at critical conditions, multidimensional detonation propagation, and detonation failure.

Algebraic Multigrid in an Industrial Environment

Dr. Klaus Stüben

Fraunhofer-Institute for Algorithms and Scientific Computing (FhG-SCAI)

Summary:

We have been working on the further development of a general framework for the application of algebraic multigrid (AMG) to systems of partial differential equations (PDEs). This framework is based on the idea of using an auxiliary matrix—called the "primary" matrix—for coarsening and/or for interpolation in the AMG context. A reasonable primary matrix describes the (pointwise) connectivity structure of some auxiliary "primary" unknown which should, in some sense, represent the connectivity structure for all "real" unknowns in the given system of PDEs and, thus, can be used for pointwise coarsening all unknowns simultaneously.

In practice, there are various ways to apply this framework to solve a given system of PDEs. For instance, if the geometric positions of the mesh points are known to AMG, the coefficients of the primary matrix may be based on geometric distances between points. Alternatively, it may be based on a suitable norm of neighboring blocks of unknowns. While, in the first case, coarsening would be closely related to geometric coarsening, it would correspond to a block-coarsening in the second case. One can also imagine that the primary matrix is defined based on some natural physical quantity for which there is no reasonable equation contained in the original system of PDEs. An example of such a situation would be the pressure in the context of the Navier–Stokes equations.

In addition, there are various ways to construct interpolation based on the resulting coarse levels. In particular, interpolation may be different for each physical unknown (e.g., based on the matrix rows in the original system), it may be the same for each unknown (e.g., based on the matrix rows of the primary matrix), or it may be blockwise.

For the framework to be as flexible as possible, a reasonable primary matrix is either defined internally to AMG or it may be user-provided by rendering the original PDE system by a primary matrix that is fully decoupled from the rest of the system. The latter option makes particular sense in situations where AMG cannot construct a reasonable primary matrix based solely on algebraic information. In many such cases, the user of AMG will be able to define a reasonable matrix based on the underlying physics of the given problem.

Projects in Computational Ecology

Tonya Kostova Vassilevska
Bulgarian Academy of Sciences

Summary:

We focus on the application of mathematical and computational methods to problems from biology. The analytical studies of small, specific ODE systems were the goals of two collaborations, which resulted in two papers on the existence of periodic solutions of some equations with applications in ecology and on the qualitative behavior and bifurcation phenomena of the FitzHugh–Nagumo system.

In addition, collaboration with the Environmental Restoration Division on the effect of oil-production sites on the ecology and biodiversity of habitats has been started. A conceptual object-oriented model has been formulated which will be implemented in C++ software. The model will use a geographic information system as a database. The data will be used to simulate the development of a foodweb consisting of vegetation, rodents, and their predators on a large-scale spatial grid in real time. Further, we shall simulate the appearance of spills and their effect on the foodweb.

A Robust Multilevel Nonlinear Method

Irah Yavneh
Technion Institute

Summary:

In a new collaboration during the summer, mainly with Carol Woodward, Panayot Vassilevski, and Jim Jones of CASC, we began investigating a novel approach for multigrid algorithms for nonlinear problems. There are two "classical" approaches for solving discretized nonlinear partial differential equations by multigrid methods. One is to perform a global linearization (GL), usually by Newton's method or some inexact variant thereof, and solve the resulting linear system approximately by a linear multigrid algorithm; this is then repeated iteratively. The second approach (FAS, by A. Brandt, or NLMG, by W. Hackbusch) is to only perform a local linearization (LL) in the error-smoothing (relaxation) process. Convergence acceleration is then provided by nonlinear coarse-grid operators.

For "nice" problems, both approaches work well and the difference in efficiency is usually not large. But for more difficult problems, the two approaches often exhibit distinct behavior, with GL holding an advantage in some situations and LL in others. We were therefore motivated to develop a method which will be at least as good as the more suitable of these two approaches, and often better than both.

Multigrid methods for nonlinear problems have been studied very extensively. For many problems, robust multigrid methods (such as "Algebraic" and "Black Box" multigrid) are well-known to be effective even in the presence of discontinuous coefficients and domain boundaries. This represents an advantage for GL methods, where the problem solved per iteration is linear. LL approaches, on the other hand, generally require "direct" rediscretization of the nonlinear operator on the coarse grids, which is less robust. However, the attraction basin of fast GL methods is usually small, which means that slow global search methods must often be applied before the GL method becomes effective. This may represent a significant disadvantage.

In a new approach, called a Multilevel Nonlinear Method (MNM), we eliminate the compromise between global and local linearization, by splitting the nonlinear operator into two parts. One is a relatively large linear part (normally obtained by Newton linearization), and the second is the small nonlinear part, which remains after the linearized operator is subtracted off. We use the non-robust FAS-like coarse-grid approximation only for the nonlinear part, while the linear part is approximated by a robust Galerkin coarse-grid operator. We thus gain the advantages of both classical approaches: a fast asymptotic convergence rate with a large attraction basin. The additional computational cost and storage is only a fraction of that of the standard algorithms, partly because the extra work is only performed on the non-finest grid. An adaptive MNM approach and a parallel version have also been conceived.

So far we have only tested the new method on one-dimensional problems, albeit some realistically difficult ones. In particular, we experimented with a nonlinear diffusion problem modeling soil water retention. The results are very encouraging, with the new method performing as predicted by our (so far simplified) analysis. Further developments, both analytical and experimental, are in progress.